



Blacklight at Pittsburgh Supercomputing Center Shines Light on Life Sciences Research

January, 2012

Abstract

Aided by a \$2.8 million award from the National Science Foundation (NSF), and offered as a unique computational resource through the Extreme Science and Engineering Discovery Environment (XSEDE) program, the installation of “Blacklight” is enabling fantastic discoveries by researchers. This system, which the Pittsburgh Supercomputing Center (PSC) acquired in July 2010, features the SGI scalable, shared memory computing platform and associated disks. The SGI® UV™ 1000 system’s extremely large, coherent shared memory opens new computational capability for U.S. scientists and engineers.

TABLE OF CONTENTS

1.0 Hardware and Software Combination Makes Running Jobs Easy	3
2.0 Blacklight Memory Advantage Program	3
3.0 Researchers from All Over Taking Advantage of the System	3
3.1 Modeling Tricky Biomolecules	4
3.2 Understanding How Gene Expression Varies	4
4.0 PSC Collaborates with SGI on Shared Memory Enhancements	5
5.0 About XSEDE	6
6.0 About the Pittsburgh Supercomputing Center	6

1.0 Hardware and Software Combination Makes Running Jobs Easy

Blacklight is an SGI UV 1000 shared memory system comprising 256 blades. Each blade holds two Intel® Xeon® X7560 eight-core processors, for a total of 4,096 cores across the entire machine. Each core has a clock rate of 2.27 GHz, supports two hardware threads and can perform nine Gflops. Thus, the total floating point capability of the machine is 37 Tflops. The 16 cores on each blade share 128GB of local memory, giving each core 8GB of memory with a total machine capacity of 32TB, though the machine is also partitioned such that it can look like two 16TB machines. Blacklight runs on SUSE® Linux® Enterprise Server. Enhancements to the operating system that were necessary to enable a machine like Blacklight have been made available to the community so they may benefit from SGI and PSC's shared experience.

The Intel® C, C++ and Fortran compilers and the GNU Fortran, C and C++ compilers are installed on the system, as are the facilities to enable researchers to run threaded, MPI and hybrid threaded programs. OpenMP programs are commonly run on the system. UPC and Java are also available on the machine. Various performance tools, debuggers and libraries are available, as well.

2.0 Blacklight Memory Advantage Program¹

“Shared memory” means that a system’s memory can be directly accessed from all of its processors, as opposed to distributed memory, where each node’s memory is directly accessed only by the processors on that node. Because all processors can share and access a single copy of data, a shared memory system is generally easy to program and use.

To help researchers take advantage of Blacklight, PSC provides a Memory Advantage Program to develop applications that can effectively use Blacklight’s shared memory capabilities. These include rapid expression of algorithms, such as graph-theoretical software, for which distributed memory often presents obstacles, and interactive analysis of large data sets, which often can be loaded in their entirety into Blacklight’s shared memory. For such projects, a PSC consultant can provide advice on debugging, performance-analysis and optimizations.

3.0 Researchers from All Over Taking Advantage of the System

“Because of the extraordinary memory size and relative ease of programming made possible by this system’s shared memory structure, scientists and engineers are able to solve problems that were heretofore intractable,” says PSC scientific directors Michael Levine and Ralph Roskies in a joint statement. “For many research communities, including data analysis and many areas of computer science, it has opened the door to the use of high-performance computation, thereby expanding the abilities of scientists to ask questions and gain insight.”

As of early 2012, a total of 1,316 users are using Blacklight for 373 projects in machine learning, natural language processing, software security, analysis of molecular biology simulations, extreme-scale performance engineering, chemistry, fluid dynamics, the early universe, condensed matter, seismic analysis, nanomaterials, astrophysics, geophysics, climate modeling and genomics.

“Researchers are extremely excited about what they are achieving with 16 terabytes of coherent shared memory, or even accessing all 32 terabytes as can be accomplished with a bit more programming,” says Nick Nystrom, PSC’s director of strategic applications. “Also revolutionary is that Blacklight offers unparalleled ease of use for rapidly testing new ideas and for expressing complex algorithms, dramatically increasing users’ productivity.”

¹ http://psc.edu/publicinfo/news/2011/110911_BLPProgress.php

3.1 Memory Scaling

Niel Henriksen of the Thomas Cheatham Research Group at the University of Utah has been running jobs on Blacklight that use more than 2TB of memory. Henriksen said, “I’m really impressed with how fast these calculations can run on Blacklight. I’ve previously done these on large distributed memory machines and, needless to say, it takes forever.”

In order to reparameterize molecular dynamics (MD) force fields for RNA, a notoriously tricky biomolecule to model, Henriksen is doing large, diffuse, polar basis set quantum mechanical (QM) calculations on mononucleotides with about 30 atoms. The jobs that use large memory are RI-MP2 calculations. Henriksen describes why they need Blacklight’s large memory and gives some very interesting initial results by stating, “Our group mostly uses MD simulations to study biomolecular structure and function. We focus on nucleic acids. We have noticed that conventional MM force fields do not properly represent the sugar pucker conformations in some regions of RNA. This causes structural pathologies to develop during simulations and limits the overall usefulness of MD simulations. To assess this problem we are using Blacklight resources to compute high level QM energy points for the pseudorotation pucker cycle of the four RNA nucleosides.”

He continues, “We are using NWChem for RI-MP2 calculations using aug-cc-pvdz and aug-cc-pvtz basis sets with the COSMO solvation model. Our preliminary results show that the MM energies correspond reasonably well with the *in vacuo* QM results. However, with the inclusion of COSMO, the pseudorotation energy profile changes dramatically. We think that if we reparameterized the MM force field to match the QM-COSMO profile, it will fix some of the sugar pucker problems we see in MD simulations.”

“Blacklight is a crucial resource for this project because its capabilities allow us to use larger basis sets and generate more accurate results. The same calculations on [the previously available system] take an order of magnitude longer and they irritate the sysadmins because the calculations swamp the disk resources. The aug-cc-pvtz calculations were not really possible [on the other system], as memory requirements are huge for aug-cc-pvtz. This project will eventually require calculations on dinucleotides as part of an effort to simultaneously optimize several force field parameters at once, thus eliminating the ‘fix one problem, cause another’ situation we have seen,” Henriksen concludes. With the current computation resources offered by Blacklight, the QM calculations on systems twice the size that are currently being run are still possible.

The reparameterization approach has been used in the past to parameterize the chi torsion angle for RNA (for instance, see pubs.acs.org/doi/full/10.1021/ct200162x on which Cheatham was co-author). This past work did not use Blacklight, but it shows that the approach works well. Sugar pucker is more complex than the single chi torsion which makes it critical to use the larger calculations possible with Blacklight to achieve the required high level of accuracy.

3.2 Understanding How Gene Expression Varies

Matthew MacManes of the Q3B Institute and Molecular & Cell Biology Department at University of California-Berkeley, states, “Blacklight is a key resource for my analyses of next generation sequence data. Without it, I would simply be unable to complete the requisite analyses. I feel so strongly about Blacklight that I have referred colleagues and collaborators. Currently, there is simply no better resource for doing this type of work.”

MacManes states, “Next generation sequencing promises to revolutionize the way in which biological phenomenon are studied. Unlike traditional Sanger sequencing, where production was measured in hundreds or thousands of bases, next generation sequencing, specifically Illumina sequencing, can produce billions of base pairs of data every day from whole genomes, specific tissues, or parts

of the genome, during specific stages in an organism's development. While this newfound ability to generate massive quantities of data has opened up novel avenues of study, it has challenged the existing computational infrastructure. Specifically, we are using next generation sequencing to better understand the genetic underpinnings of complex phenotypes in non-model mammals. In particular, I am interested in how differences in gene expression are related to the transition from solitary to social living in the Colonial tuco-tuco (*Ctenomys sociabilis*).

He continues, "Using RNA extracted from the hippocampus of captive animals housed in either social or solitary conditions, we have begun to understand how gene expression varies. While these studies are still in their infancy, several genes not previously implicated in social behavior have been revealed. Additionally, analyses of natural selection, nucleotide diversity and gene pathways/ontology are in progress."

"Because hundreds of gigabytes of raw data are used in a single analysis, the shared memory architecture of Blacklight is essential for my analyses. In fact, I have used, unsuccessfully, several other XSEDE resources. Analyses require between 70GB and 700GB of RAM, depending on the amount of data and complexity of the analyses. To the best of my knowledge, Blacklight is the only resource on which these analyses are possible. In addition to hardware, Blacklight has an extensive collection of pre-compiled modules for the analyses of next-generation sequencer data. Because of the complex nature of all XSEDE resources, along with (typically) the multiple dependencies of the programs used for analyses, compiling the 20 programs I typically use is very time consuming, often requiring multiple interactions with tech support. Having these programs installed and maintained by PSC staff is extremely helpful," MacManes concludes

4.0 PSC Collaborates with SGI on Shared Memory Enhancements

As a resource of XSEDE, the NSF cyber-infrastructure program, Blacklight serves researchers nationwide across many fields of science and engineering. To optimize Blacklight's productivity for these researchers, PSC staff, under the lead of J. Ray Scott, Director of Systems & Operations, worked closely with SGI to enhance a number of Blacklight features. These enhancements include:

- Substantially reducing the boot time of the system
- Enabling high input-output file performance for very large datasets without the need to change application software
- Developing new tools to diagnose research application problems — to optimize job layout among processors and to track and isolate hardware problems
- Modifying the portable batch system scheduler so that resources assigned to one shared memory job don't interfere with other jobs
- Enabling researchers to use Zest, a high-speed memory snapshot utility, and SLASH2, a wide-area data-replicating file system, to ease the burdens of data management for very large, distributed datasets

According to PSC and SGI officials, many of these enhancements will benefit not only XSEDE and the researchers who use Blacklight, but also the research community at large, by enabling improved performance for all large-configuration SGI UV 1000 systems. "We've been at this for 25 years," said PSC scientific directors Roskies and Levine, "and we've become expert at working out the sometimes perplexing problems that go along with bringing a new system through the process of shakedown and final testing to provide researchers with a maximally productive tool."

5.0 About XSEDE

The Extreme Science and Engineering Discovery Environment (XSEDE) is the most advanced, powerful and robust collection of integrated advanced digital resources and services in the world. The five-year, \$121-million project is supported by the National Science Foundation, and replaces and expands on the NSF TeraGrid project. XSEDE lowers technological barriers to the access and use of computing resources. Using XSEDE, researchers can establish private, secure environments that have all the resources, services and collaboration support they need to be productive. The XSEDE partnership includes over a dozen universities and research centers, and is led by the University of Illinois's National Center for Supercomputing Applications (NCSA).

6.0 About the Pittsburgh Supercomputing Center

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. Established in 1986, PSC is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry, and is a partner in the National Science Foundation XSEDE program.

Global Sales and Support: sgi.com/global

©2012 Silicon Graphics International Corp. All rights reserved. SGI, the SGI logo and UV are registered trademarks or trademarks of Silicon Graphics International Corp. Intel, Xeon and the Intel Xeon logo are registered trademarks of Intel Corporation. All other company and product names are registered trademarks or trademarks of their respective companies. 05032012 4347

